



"Method and System for Bandwidth Estimation"

Introduction

- 5 The present invention relates to a method and system for estimating the effective bandwidth on-line at a node in a communications network.

Due to the recent explosion of Internet use, there has been an enormous increase in data or traffic flow being sent over communication networks. This has led to a
10 huge increase in bandwidth requirements by operators which is expensive. A problem with handling traffic flow is that it is difficult to estimate due to the "bursty" nature of traffic flows. For example, if someone decides to send video data to another source over a communications network, it uses a lot of bandwidth. A solution to this problem is that when network operators are allocating bandwidth to
15 a node or a router, a large amount of bandwidth is allocated to that node to handle any of the bursty traffic. This is very wasteful of bandwidth.

Various methods and systems have been employed to control the admission of traffic flows in packet switch networks such as Internet Protocol (IP) based
20 networks which includes traditional IP, IntServ, DiffServ and MPLS networks as well as ATM and Frame Relay. It is important to have a reliable estimation of resources required to meet quality of service performance requirements. The estimation of resources can be based on the usage of measurement and estimation of effective bandwidth. At present it is impractical to estimate the
25 required resource for each traffic flow separately on-line for a number of reasons, including scalability, as well as the fact that many traffic flows exist only for a short period of time, especially in IP based networks. It is more practical to estimate effective bandwidth for traffic aggregates comprising many flows.

- 30 The effective bandwidth of a variable bit rate traffic flow is a measure that summarises statistical properties of the traffic flow into just one value, while taking into account the quality of service requirements such as cell loss ratio, delay, and other parameters. Also, the bandwidth estimation of the effective bandwidth can take into account some other parameters such as link, bandwidth, buffer size, etc.

On-line implementation of the estimation of the effective bandwidth is complicated due to a number of reasons, such as performance related constraints, measurements and estimation, technique constraints, router constraints and measurement point constraints. The affect of all these factors is that on-line
5 estimation of effective bandwidth at a node has not been achieved accurately.

The estimation of the effective bandwidth of a traffic flow consumes router resources. Therefore, it is impractical to estimate the effective bandwidth for each traffic flow when the number of flows becomes too large. Such a straight forward
10 approach leads to a scalability problem as the number of flows increase. Some models used in the estimation of the effective bandwidth rely on parameters that are either difficult to measure effectively or reliably. An example of such a parameter is the peak rate of flow. The peak rate of flow is very sensitive to local fluctuations of traffic flow. It's estimation is often based on averaging over a small
15 period of time and it is either over estimated or under estimated. It has been found that the declaration of peak rate is not accurate thus it does not provide a full remedy to the problem.

Router constraints is a problem because the buffer counters within the router
20 provide raw information for estimation of effective bandwidth only. Some routers provide access to counters at output buffers but for bandwidth estimation purposes, it is not an accurate estimation. This is also applicable to other types of network elements including switches or bridges.

25 The mean rate of traffic flow is a good characteristic for bandwidth estimation. Mean rate is a very robust characteristic of traffic flow. In particular, it is not practically sensitive to local short term fluctuations in flow rate. Mean rate is relatively simple to measure. For this reason, it is very desirable to use the mean rate as the basis for estimation of effective bandwidth. Unfortunately, in general
30 the mean rate of the source does not reflect a network resource required for a flow to meet quality of service requirements. It has also been found that the direct usage of mean rate, together, for example, with the peak rate for the estimation of the effective bandwidth can lead to inaccurate results. A principal drawback of such a method or system, as already mentioned, is that the peak rate is very

- 3 -

sensitive to local fluctuations of traffic and is difficult to measure.

5 In the paper F. P. Kelly "Effective Bandwidths at Multi-class Queues", Queuing Systems 9, 1991, pp 5-16" a method for the estimation of the effective bandwidth of an individual flow is proposed. In principle this method can be used for the off-line estimation of the effective bandwidth required. For on-line usage however it is too time consuming. The other disadvantage of this method is that it does not take into account buffer size and as a consequence of this it does not deal with any delays. The paper does not propose a method for estimating effective bandwidth on-line accurately.

10

In another paper D.D. Botvich & N.G. Duffield "Large Deviations, Economies of Scale, and the Shape of the Loss Curve in Large Multiplexers" Queuing Systems 20, 293-320 (1995) investigates the behaviour of the effective bandwidth in the situation when the number of traffic sources is increased. The economy of scale in this case is evaluated and estimated for properly chosen scaling. This paper provides a theoretical basis for estimating the effective bandwidth for a number of flows off-line.

15

20 PCT Patent Application No. PCT/IE98/00013 "Telia Research AB et al" discloses a general method for the estimation of the effective bandwidth applicable both to individual flows and traffic aggregates. This method is flexible and accurate for the estimation of effective bandwidth for individual flows. This method can be used both off-line and on-line. The on-line implementation however has some limitations due to the fact that most of today's routers and switches are not able to provide sufficiently accurate "raw" measurement data. Thus, the on-line estimation of effective bandwidth is invariably inaccurate for a large number of flows.

25

30 In this specification, the term "traffic flow" is used interchangeably with other terms including "call" and "connection". It will be appreciated to someone skilled in the art that the term "traffic flow" is used in IP based networks, while the terms "call" and "connection" are used in ATM networks. For this reason we also use terms "packet" and "cell" interchangeably. Further, in this specification, we use the terms "measurement" and "estimation" in a different way. We refer to measurement as

the process of collecting the raw traffic information. The estimation process uses the data collected from the measurement process.

5 In view of the above problems, there is a need for a system and method for estimating the effective bandwidth at a node in a communications network. The bandwidth estimation must be quality of service sensitive and deal with loss and delay parameters in real time. Further, the object of the present invention is to provide a solution to the above problems and provide a bandwidth estimation technique in real time that relies on robust and easily measurable parameters in the
10 network.

Statements of Invention

According to the present invention there is provided a method of estimating the
15 effective bandwidth at a node in a communications network comprising:-

identifying a type of traffic flow off-line;

20 measuring an effective bandwidth and a mean rate of traffic flow of the identified type of traffic flow;

defining a relationship between the effective bandwidth and the mean rate of traffic flow for the identified type of traffic flow;

25 storing a database of relationships for different identifiable types of traffic flows off-line;

identifying a type of traffic flow on-line at the node;

30 measuring the mean rate of traffic flow of the identified type of traffic flow on-line;

comparing the identified type of traffic flow on-line to a similar identified type of traffic flow off-line to obtain a relationship from the database of

- 5 -

relationships;

5 estimating an effective bandwidth of the identified type of traffic flow on-line at the node from the mean rate of traffic flow measured from of the identified traffic flow on-line and the relationship obtained from the similar identified type of traffic flow off-line.

10 It will be appreciated that in the steps of identifying the traffic flow, measuring parameters, defining a relationship and storing a database of relationships can be carried out at any node in the communications network. These steps are not restricted to a particular node and can be carried out at a node different to that being estimated.

15 In another embodiment, the method comprises the further steps of:-

 measuring the mean rate of traffic flow individually on-line for a number of identifiable types of traffic flows forming a traffic aggregate;

20 individually comparing each type of identified traffic flow on-line to a similar type of traffic flow off-line to obtain a relationship from the database of relationships for each type of identified traffic flow on-line;

 estimating the effective bandwidth on-line at the node for the number of identifiable types of traffic flow.

25 Further the method comprises the step of:-

 identifying a traffic aggregate on-line which represents traffic produced by a group of flows of the same type.

30 In another embodiment method further comprises the step of:-

 updating the database of relationships with a relationship from the estimated effective bandwidth and the mean rate on-line value.

Ideally the step of obtaining the effective bandwidth off-line takes account of quality of service requirements in the communications network.

- 5 In one embodiment the steps of obtaining the effective bandwidth off-line incorporates a loss ratio or delay parameter in the communications network.

Preferably the step of calculating the relationship between the effective bandwidth and the mean rate of traffic flow is calculated off-line over a number of time intervals.

10

Ideally the step of measuring the mean rate of traffic flow comprises measuring the number of packets of information at the node on-line.

15

In a preferred embodiment the relationship between the effective bandwidth and the mean rate of traffic flow is calculated as the ratio of the effective bandwidth with respect to the mean rate of the traffic flow.

20

Ideally the effective bandwidth as claimed in any preceding claim, comprising estimating the effective bandwidth off-line and measuring the mean rate off-line for a number of flows to obtain a relationship represented by:-

$$C_i = \frac{E_i}{m_i} \quad i = 1 \dots N$$

25 where

C_i = ratio for the i-th flow

E_i = Effective bandwidth the for i-th flow measured off-line

30

m_i = Mean rate of traffic flow for the i-th flow measured off-line

N = Number of different traffic flows.

- 7 -

Preferably the method comprises estimating the effective bandwidth as claimed in claim 10, comprising estimating the ratio C for a number of traffic flows off-line represented by:-

$$5 \quad C = (C_1 + \dots + C_N) / N$$

In another embodiment the method comprises estimating the typical ratio C off-line for a number of traffic flows each flow of a particular time duration represented by:

$$10 \quad C = (T_1 C_1 + \dots + T_N C_N) / (T_1 + \dots + T_N)$$

where T_i = time duration of a particular flow

N = total number of flows.

15 Preferably the method comprises estimating the typical ratio C off-line for a number of traffic flows each traffic flow of a particular time duration and a measured value of mean rate and effective bandwidth represented by:-

$$20 \quad C = (T_1 E_1 + \dots + T_N E_N) / (T_1 m_1 + \dots + T_N m_N).$$

where T_i = time duration of a traffic flow

E_i = estimated effective bandwidth of a traffic flow

m_i = measured mean rate of a traffic flow

N_i = number of traffic flows.

25 In another embodiment the method of estimating the effective bandwidth comprises the further step of:-

30 factoring the quality of service requirements into the estimation of bandwidth by calculating the maximum number of flows to maintain quality of service requirements by:-

$$\text{Prob} \{mC_1 + \dots + mC_N > B\} < P$$

- 8 -

wherein

- m = measured mean rate of traffic flow
- C = typical ratio for one traffic flow
- B = bandwidth requirement to maintain quality of service
- P = loss ratio
- N = Number of traffic flows

In another embodiment the method of estimating the effective bandwidth comprises estimating a ratio D off-line for a number of traffic flows represented by:

$$D = \frac{B}{mn}$$

where

- B = bandwidth requirement to maintain quality of service
- m = mean rate
- n = maximum number of flows allowed at the node.

In another embodiment of the present invention there is provided a system for estimating the effective bandwidth at a node in a communications network having a router associated with each node and computer hardware for monitoring traffic flow at each node comprising:-

means for identifying a type of traffic flow off-line;

characterised in that the system further comprises:-

means for measuring an effective bandwidth and a mean rate of traffic flow of the identified type of traffic flow off-line, calculating the ratio of the effective bandwidth with respect to the mean rate of the traffic flow off-line to obtain a relationship;

means for storing a database of relationships for different identifiable types of traffic flows off-line;

means for identifying a type of traffic aggregate on-line at the node and measuring the mean rate of traffic aggregate of the identified type of traffic aggregate on-line;

5 means for comparing the identified type of traffic aggregate on-line with the database of relationships to obtain a relationship representing the closest comparison to traffic flow off-line;

10 means for estimating the effective bandwidth of the identified type of traffic aggregate on-line at the node.

15 Ideally the computer hardware comprises means to measure the mean rate of traffic flow individually on-line for a number of identifiable types of traffic flows forming a traffic aggregate.

In one embodiment the computer hardware comprises means to estimate the effective bandwidth on-line at the node for the number of identifiable types of traffic flow.

20 Preferably the database of relationships is updated with a relationship from the estimated effective bandwidth and the mean rate effective bandwidth on-line value for future reference.

25 In one embodiment the computer hardware takes account of quality of service requirements when estimating the effective bandwidth at the node.

Ideally the computer hardware incorporates a loss ratio or delay parameter when estimating the effective bandwidth at the node.

30 Preferably the computer hardware calculates the relationship between the effective bandwidth and the mean rate of traffic flow calculated off-line over a number of time intervals.

In another embodiment the computer hardware comprises:-

- 10 -

a processor;

a memory storing programmed instructions; and

5

a bandwidth estimation server having an admission control unit and storing the database of relationships.

10 In a further embodiment the computer hardware is remote from the nodes in the system.

Ideally there is provided a user interface displaying estimated bandwidths on-line for different nodes of the system.

15 Ideally the admission control unit can allocate bandwidth to certain nodes within the system in response to the displayed estimated bandwidths on-line at the nodes.

Detailed Description of the Invention

20 The present invention will be more clearly understood from the following description of some embodiments thereof, given by way of example only, with reference to the accompanying drawings, in which:-

Fig. 1 is a physical representation of how the invention is carried out,

25

Fig. 2 is a flowchart of the operation of the invention.

Referring now to the drawings and initially to Fig. 1, there is provided a communications network, indicated generally by the reference numeral 1 in which a
30 number of routers 2 positioned at each node control traffic flows within the communications network 1. A bandwidth estimation server 3 communicates via the network 1 is connected to each of the routers within the communications network. The bandwidth estimation server 3 can be remote from each router 2. It will be appreciated that the bandwidth estimation server 3 can be housed directly on a router

- 11 -

2 if necessary. A computer 4 is connected to the bandwidth estimation server 3 having a user interface 5 which allows an operator see the amount of bandwidth being used at each router 2 in real time. The bandwidth estimation server 3 stores a database of relationships, which will be discussed in more detail below. The server 3 has the ability to provide the configuration information either by downloading to routers 2 or on a request from a router.

Referring now to Fig. 2, there is illustrated a flow diagram of the operation of the invention. In step 11, a type of traffic flow is identified off-line. This may be a recorded traffic flow at a previous time. In step 12, the effective bandwidth and a mean rate of traffic flow of the identified type of traffic flow is estimated off-line. In step 13, the relationship is defined between the effective bandwidth and the mean rate of traffic flow for the identified type of traffic flow. In step 14, a ratio estimate is obtained off-line. Steps 11 to 14 are carried out for a number of different identified traffic flows of the same type or as well as for different identified types of traffic flow if necessary. Each time a relationship is obtained, it is added to a database of relationships in step 15 which stores relationships for different identified types of traffic flow off-line. Steps 11 to 15 can be referred to as the tuning stage of the method. The tuning stage may be carried out on-line. In step 16, a traffic flow on-line at a node or router 2 is identified. In step 17, the mean rate of traffic aggregates the on-line measurement. In step 18, the identified on-line traffic type is compared with the database of relationships. A relationship is obtained which represents the closest off-line traffic type with the identified on-line traffic aggregate. In step 19, an estimate of the effective bandwidth on-line of the identified traffic aggregate on-line is obtained from the mean rate of traffic flow of the identified traffic aggregate on-line and the relationship obtained from the same or similar identified type of traffic off-line from the database of relationships.

Referring now in detail to the idea of the method, suppose that the total mean rate of a traffic flow or group of traffic flows is measured and is equal to m . Then the effective bandwidth is estimated to be equal to $D \times m$, where D is some constant coefficient that has been estimated during the tuning stage of the method.

It is possible to estimate the effective bandwidth by measuring only mean rate m of

the traffic flow. It should be emphasised that the coefficient D is always known as it is defined during the tuning stage which will be discussed later. In general, it is useful to think about D as a "typical ratio" of effective bandwidth with respect to the mean rate of traffic flow or traffic aggregates (i.e. group of related flows).

5

The basic idea of the method can be illustrated as follows. A parameter D is estimated in the following way. Suppose that there is provided a collection of N recorded flows with measured mean bit rates m_1, \dots, m_N , respectively. Each traffic flow has a flow duration denoted by T_1, \dots, T_N , respectively. An algorithm is used off-line to estimate effective bandwidths of flows. E_1, \dots, E_N are the corresponding effective bandwidths of each particular flow. It is implicitly assumed that quality of service requirements such as loss ratio and delay as well as other parameters such as link bandwidth and buffer size are used for the estimation of E_1, \dots, E_N . The estimation of a typical ratio of the effective bandwidth with respect to mean rate for a traffic flow. In principle, there are a few ways to do this, depending on which parameters are important.

Put

$$C_i = E_i / m_i, \quad i = 1, \dots, N.$$

Where C_i = typical ratio for the i-th flow

E_i = effective bandwidth of the i-th flow

m_i = mean rate of the i-th flow

25

Then C for a traffic flow or number of traffic flows can be estimated as

$$C = (C_1 + \dots + C_N) / N. \quad (1)$$

30

The method simply puts $D = C$. With a flow or a traffic aggregate with mean rate m , it is possible to interpret $D \times m$ as the estimation of the effective bandwidth for the flow or a traffic aggregate on-line. For one flow or a traffic aggregate comprising a small number of flows, such estimation is not very accurate but for a traffic aggregate

comprising many flows, the estimation is accurate. This case is the most simple one as we do not know how long the duration of traffic flow. The value of D can be used to relate the mean and the effective bandwidth of a large number of flows. This gives the possibility to express flow bandwidth requirements on-line in a very simple way, i.e. $D \times m$, where D is some constant coefficient that has been estimated during the tuning stage of the method and m is the mean rate measured on-line of the traffic flow. The value of D is obtained from the database of relationships of off-line recorded traffic flows. The value D that is used is the closest compared traffic flow off-line to that of the on-line traffic flow. This is the basic idea of the invention.

The expression (1) for "typical ratio" does not take into account flow duration, absolute value of mean rate but only the relative relationships of effective bandwidth and mean rate. The absolute values of mean rates and effective bandwidths can also be useful in bandwidth estimation. The described statistical procedure above can be too simplistic in real network settings. The following formula can be more accurate for the estimation of C in practice:-

$$C = (T_1 C_1 + \dots + T_N C_N) / (T_1 + \dots + T_N). \quad (2)$$

$$C = (T_1 E_1 + \dots + T_N E_N) / (T_1 m_1 + \dots + T_N m_N). \quad (3)$$

Formulae (1), (2) and (3) reflect different approaches to how to construct a sampling process. The use of the formula depends on the problem and the availability of the information for the estimation of C such as flow durations T_1, \dots, T_N . Usually the more parameters that are available for the calculation of the typical ratio the more accurate the bandwidth estimation will be. The accuracy of the effective bandwidth estimation can be influenced by different factors. For example, it has been found in particular that the estimation of effective bandwidth for individual flow has to be robust. Also the flows should have the same, or at least, close enough quality of service requirements.

In experiments, it has been found that the effective bandwidth estimation produces some errors for the actual values of effective bandwidth of individual flows. For some flows, the estimated value will be less than the actual one and, for others,

more. What is important is that the total requirements for the bandwidth will be exactly the same in both cases, based on the estimation of effective bandwidth of individual flows or by the above described method.

5 The actual estimation of C can be carried out in a few different ways. For this, it is required to estimate the mean rate, the effective bandwidth and time duration of a number of typical flows. The following are ways to do this:

1. Off-line estimation for "typical" recorded sources;
- 10 2. On-line selective estimation for some sources;
3. On-line estimation for all or almost all sources during some time period.

Note that the third variant does not assume constant measurements of effective bandwidth. The third variant gives better approximation for typical ratios but can
15 have similar disadvantages as on-line implementation of the effective bandwidth. The first variant is the most simple for the implementation, but is less accurate than the others.

It will be appreciated that what has been described is a simple model to carry out
20 the method. We now consider more advanced mathematical models of the method to estimate the effective bandwidth. In the models, quality of service parameters are taken into account by using large deviation theory to estimate the parameter D or C.

25 A few models are described below. The models include the so-called homogeneous case model and models 1, 2 and 3. Homogeneous case model is a specific one designed to deal with the case when mean rate is the same for all traffic flows. This model can be easily adopted to be practically implemented in a communications network. The other three models 1, 2 and 3 are generic and more
30 complex. Model 1 is a generalization of the homogeneous case. Models 1, 2 and 3 are different in terms of accuracy of the estimation of the effective bandwidth as well as in the implementation of complexity. It should be noted that when implementing the proposed method, it is required to find proper tradeoffs between the required accuracy and the implementation complexity at each node or router 2

in the communications network 1.

Homogeneous Case Model

- 5 The above method is accurate when ratios C_i of effective bandwidths to mean rates would be the same for all flows or close enough to mean typical ratio C . Obviously, when the variance of ratios C_i increases, the uncertainty about the resource required by each flow also increases. As a result, the communications network has to allocate more resource for each flow in order to meet the quality of service
- 10 requirements. It is important to take this factor into account. Suppose that flows share a link of bandwidth B . Suppose that all flows have the same (constant) mean rates m , but all C_i are random and have the same distribution function F . Distribution function F is estimated as empirical distribution function for C_1, \dots, C_N . Also, suppose that all flows have loss ratio equal to p . It is important to find the
- 15 maximum number of traffic flows that can be admitted to the router 2 while ensuring quality of service requirements are met, represented by " n ".

The resource will depend on how many traffic flows the network could carry out while still meeting the quality of service requirements. In statistical terms, this

20 condition can be expressed as follows. It is necessary to find the maximum value n such that:-

$$\text{Prob} \{ m C_1 + \dots + m C_n > B \} < p. \quad (5)$$

- 25 In other words, it is necessary to find the maximum value n such that the saturation probability, that the total resource $m \times (C_1 + \dots + C_n)$ that is required by all flows is greater than B bandwidth requirement to maintain quality of service requirements, is less than cell loss ratio p . When the maximum number n of flows is known, then the value of D can be estimated as:-

30

$$D = B/(m n).$$

It is possible to obtain n by some approximation methods, for example, based on large deviation theory.

Model 1

Model 1 is a direct generalisation of the homogeneous case. In this case, we take
 5 m equal to average mean rate estimated from different samples:

$$m = (m_1 + \dots + m_N) / N.$$

The rest of the calculations are the same as in the homogenous method.

10

Model 2

Again, it is assumed here that flows share a link of bandwidth B. But suppose that
 flows can belong to one of the following types: each flow of the j-th type has
 15 estimated mean rate m^0 . Note that it is not assumed here that the mean rates of
 flows are constant. The ratio of effective bandwidth with respect to mean rate C_{ji} for
 the i-th flow of type j is a random variable with some distribution function F^0 , $j=1,$
 ..., J, depending on type of flow. Suppose also that all flows have cell loss ratio
 equal to p.

20

The problem of finding n_j can be done in a similar way as in the homogeneous case
 dealing with each class j separately in order to find D_j .

The values of D_j can be used separately if aggregate measurements for each traffic
 25 flow are available. But in the case when such measurements are not available then
 it could be possible to estimate value of p_j presenting the fraction of traffic class j,
 then typical ratio D for all classes is equal to

$$D = p_1 D_1 + \dots + p_J D_J$$

30

Model 3

The previous model can be generalised even further. Assume that flows share a
 link of bandwidth B. Assume that random vector (m_i, c_i) associated with each flow

has a distribution function $F(\dots)$, where m_i and $E_i = m_i \times C_i$ are mean rate and effective bandwidth of the flow, respectively. It is assumed that effective bandwidth is calculated according to some specified cell loss ratio p or delay d . This model is a bit more complicated but still can be treated in similar way as Model 2.

5

It will be appreciated that the method is based on the observation that the ratio of effective bandwidth with respect to the mean rate does not vary. Tuning of the method parameters can be formed off-line using recorded traffic flows or on-line using real time traffic flows. The actual usage of the method is very simple. It relies on the measurement of mean rate of traffic flow or traffic aggregates. The effectively bandwidth of traffic aggregate is simply proportional to the mean rate $D \times m$ where m is measured mean rate of traffic and coefficient D is estimated during the tuning phase of the method. Steps 11 to 15 are the tuning phase of the method.

15

It will be appreciated that the invention works equally well for a single traffic flow or a group of traffic flows or traffic aggregates. It will also be appreciated that after an on-line estimation of bandwidth, the database of relationships can be updated with the estimated values.

20

It has been found to improve the accuracy of the estimation of effective bandwidth off-line, the mean rate of traffic flow is calculated over a number of time intervals. It will be appreciated that in this specification, when we are discussing measuring the mean rate of traffic flow, it corresponds to the measured number of packets of information at the node or router on-line.

25

It will be further appreciated that a network operator controlling the communications network can see from the user interface 5, the bandwidth requirements and the number of traffic flows going through each router 2 in the network. If the number of traffic flows going through the router 2 is too large, the invention estimates the bandwidth in real time so that the operator can see at a glance that too much traffic is being directed through that router. An admission control unit located within the bandwidth estimation server 3 can take action by allocating more bandwidth to that router 2 or rejecting some of the calls going through by re-routing the calls via

30

- 18 -

another source. Thus, call admission control is achieved in real time without any loss of quality in the communications network.

It will be further appreciated that the mathematical models described in this specification are just some techniques and it is envisaged that other models may be used using the same parameters to estimate the effective bandwidth by obtaining the typical ratio.

It will be appreciated that the traffic aggregate is an aggregation of traffic flows going via a router. In general, the traffic aggregate can be defined via a filter that "selects" packets based on the content of packet headers or other packet data, and/or on implicit or derived attributes associated with the packet. The filter is a set of conditions on the components of a packet's classification key. A filter is said to match only if each condition is satisfied. Some examples of traffic aggregates are traffic that goes via a link in one direction, traffic of the same class that goes via a link in one direction (DiffServ), traffic of the same label switched path or the same forwarding equivalence class that goes via a link in one direction (MPLS), traffic with the same destination address or group of addresses, traffic of the same source and destination addresses, and traffic of the same application or port number.

It will be appreciated that various aspects of the invention may be embodied on the computer 4 that is running a program or program segments originating from a computer readable or usable medium, such medium including but not limited to magnetic storage media (e.g. ROMs, floppy disks, hard disks, etc.), optically readable media (e.g. CD-ROMs, DVDs, etc.) and carrier waves (e.g., transmissions over the Internet). A functional program, code and code segments, used to implement the present invention can be derived by a skilled computer programmer from the description of the invention contained herein.

In the specification the terms "comprise, comprises, comprised and comprising" or any variation thereof and the terms "include, includes, included and including" or any variation thereof are considered to be totally interchangeable and they should all be afforded the widest possible interpretation.

The invention is not limited to the embodiments hereinbefore described but may be varied in both construction and detail.